

SUMEX-AIM PROSPECTUS

INTRODUCTION AND OBJECTIVES

In partnership with and with financial support from the Biotechnology Resources Branch (BRB) of the Division of Research Resources, National Institutes of Health (NIH), Stanford University is developing and operating a National Shared Computing Resource to explore advanced applications of computer science in health research. The SUMEX (Stanford University Medical Experimental Computer) facility, under Professor Joshua Lederberg as Principal Investigator, is national in scope in that a major part of its computing capacity will be made available to authorized research groups throughout the country by means of a communications network. There are two main objectives of the facility: 1) the specific encouragement of applications of artificial intelligence in medicine (AIM) and 2) the managerial, administrative, and technical demonstration of a national shared technological resource for health research.

The emergence of more economical technologies for data communications allows a liberation from geographically organized facilities in favor of the concentration of functionally specialized capabilities at different sites. In addition to the economic advantages of resource sharing, a closer interaction between diverse research efforts is expected to promote a more systematic exchange of research products and ideas. This is particularly true in computer science where technical complexities have tended to encourage the development of relatively isolated groups, each pursuing a line of research and program development and making limited use of the working programs available from others. The SUMEX-AIM project seeks to lower these barriers in the specific area of artificial intelligence applied to health research. Indeed, multilateral community-building rather than unilateral service is the project's essential mandate.

"ARTIFICIAL INTELLIGENCE" RESEARCH

The term "artificial intelligence" (AI) is applied to research efforts aimed at studying and mechanizing information processing tasks that have generally been considered to require human intelligence. The current emphasis in the field is on efficient acquisition and utilization of material knowledge, and representation of conceptual abstractions in problem solving processes. AI systems are characterized by complex information processes that are, to a large extent, non-numeric, e.g., graph searching and symbolic pattern analysis. They involve procedures whose execution is controlled by different types and forms of knowledge about a given task domain, such as models, and fragments of "advice" in the form of systems of constraints or heuristic rules. Unlike conventional algorithms commonly based on a well-tailored method for a given task, AI procedures typically use a multiplicity of methods in a highly conditional manner - depending on the specific data in the task and on a variety of sources of relevant information.

For example, an AI system for data interpretation would assist a user at a higher conceptual level than the application of a numerical algorithm for curve fitting by suggesting and evaluating analytical expressions to match the data. Thus it may include deciding what model best describes the data source and what actions to take as a result. Such a model would be based on the data itself and on other pieces of information which may be relevant to constraining a solution. Inherent in these processes are a store of general symbolic knowledge about the problem domain describing what is fact, what is reasonable and unreasonable, and what solutions have worked in the past. Using an ability to communicate effectively with the problem data source and a human user, the programs might examine the knowledge base to construct and test plausible explanations of a particular set of data and to project and decide among subsequent alternative courses of action.

This type of "intelligent" assistance by computer program is perhaps best illustrated by a number of examples taken from on-going research efforts identified as initial users of the facility. The DENDRAL project at Stanford is aimed at assisting the biochemist in interpreting molecular structures from mass spectral and other chemical information. In cases where the characteristic spectrum of a compound is not catalogued in a library, these programs carry out the rather laborious processes a chemist must go through to interpret the spectrum from "first principles." By symbolically generating "reasonable" candidate structures from hints within the spectrum and a knowledge of organic chemistry and mass spectrometry, the program infers the unknown structure to be the one which best explains the observed spectrum. There is no direct algorithmic path available to

determine such a molecular structure from the spectral data - only the inferential process of hypothesis generation and testing within the domain of reasonable solutions defined by a knowledge of organic and physical chemistry.

This process, as implemented in the computer, is a simplified example of the cycle of inductive hypothesis - deductive verification that is often taught as a model of the scientific method. (Whether this is a faithful description of contemporary science is arguable; and how it may be implemented in the human brain is unknown. In any case, these are useful leads rather than absolute preconditions for the pragmatic improvement of mechanized intelligence for more efficient problem solving.) The elaboration of these approaches as deeply as we can with existing hardware and software technologies is the most promising approach to enhancing the application of the computer to the vaguely structured problems that dominate our task domains.

A project related to DENDRAL, carried out in collaboration with the University of California at San Diego, seeks to infer the structures of complicated proteins from x-ray crystallographic data through a similar paradigm of hypothesizing reasonable structures and testing their ability to explain the observed data.

Other projects at Stanford and Rutgers University seek to assist experimentally in diagnosing disease and suggesting treatment. In these cases the input data are clinical symptoms and physical and biochemical measurements and the knowledge domain is the physiology and pathology of various organ systems as well as the effects and interactions of courses of treatment. Current applications of this work are in the areas of infectious diseases and diseases of the eye. Long term applications of these types of computer programs might be to consolidate and reconcile the knowledge from a diverse group of experts or to enable more effective treatment of disease in locations that lack access to specialized expertise.

Other examples, in areas of psychology, are aimed at building and testing complex models of human cognitive and affective processes. Programs are being designed at Stanford to simulate paranoid or other behavior patterns in response to natural language discourse. Embedded in such programs are a general symbolic model for the behavior type and an ability for natural language communication. In response to a human user, the program seeks to understand input discourse in terms of the behavioral model and to produce appropriate English language sentences in response. Applications of this work may be to improve our understanding of particular behavior patterns by systematizing the characteristics of models which emulate them or to assist in training medical students in psychiatry.

A system is being developed at Rutgers which takes as an input a social episode, i.e., an account of a sequence of actions involving the interactions of several persons within some social context, and it generates an interpretation of the episode in terms of intentions and reasons that might have motivated the persons in the episode to perform their actions. This system is based on a model consisting of a body of rules of belief (about specific people, actions, and motivations) and a strategy of interpretation. Applications of the system may be in improving communication processes such as the psychiatric interview.

These examples are given for the purpose of concrete illustration of the types of problem areas we seek to explore. They are not intended to bound the domain but rather to stimulate new ideas along the lines of "intelligent" programs with medical applications. These terms are not precisely defined at this time and our objectives might better be phrased in terms of "advanced computer science concerned with mechanized theory formation and problem solving in medical research and practice." We will emphasize the past achievements of AI-oriented and AI-labeled research in providing facilities for further advances such as symbolic knowledge representation and manipulation, concept formation, problem solving, learning, and natural human communication (e.g., language, speech, vision, etc.),

MANAGEMENT AND USER QUALIFICATIONS

The SUMEX-AIM facility is community-oriented and its organizational structure is being established to emphasize user support and interactions among user groups. The following summarizes the overall structure as it is tentatively constituted. Additional detail will be provided as appropriate, giving specific information about how prospective users may gain access to the facility.

The user community is divided for administrative purposes into two groups: 1) those at the Stanford Medical School (local to the facility) and 2) those elsewhere in the country and at Stanford. The facility resources (computing capacity and manpower) will be allocated in equal portions to these two groups. Stanford Medical School users will be qualified for access to the facility by Dr. Lederberg in his capacity as Principal Investigator for the SUMEX grant. The national user group will gain access to and be represented in the design, development, and allocation of the facility resources through an advisory group for a national program in Artificial Intelligence in Medicine (AIM). The AIM Advisory Group will consist of members at large of the AI and medical communities, facility users, and the Principal Investigator of SUMEX as an ex-officio member. A representative

of the NIH-BRB will serve as Executive Secretary. It will advise the AIM Executive Committee, whose responsibility is to give overall direction to the national AIM program. The AIM Executive Committee consists of the Principal Investigator of the SUMEX Project, the Principal Investigator of a series of workshops on AI applications in medicine, a representative of the NIH-BRB, and a representative of the AIM Advisory Group.

Besides its charge to review potential uses of the SUMEX-AIM facility, the AIM Advisory Group will advise the AIM Executive Committee on the allocation of funds needed to assist in interfacing recommended new users. The SUMEX-AIM computing resource will be made available to qualified users without any charge, which of course entails a careful review of the merits and priorities of proposed applications. At the discretion of the advisory group, attendant communication and transportation costs to allow specific users to gain access to the facility may be covered as well.

Qualifications for new users and the details of NPAIM assistance to users will be more specific at a later time. In general terms, however, potential users will be judged on the basis of:

1. The scientific interest and merit of the proposed research,
2. The relevance of the work to the AI approach of SUMEX-AIM as may be indicated in part by the need for the specialized SUMEX facilities as opposed to other computing alternatives,
3. The prospective contributions and role of a user in the community, e.g., developing and sharing new systems or applications programs, sharing use of special hardware, etc.
4. The user's capability and intentions of operating in a community-effective style for mutual advantage. Besides the programming innovations that some users may be able to contribute, all are expected to furnish expert knowledge and advice about the existing art in the fields relevant to their special interests.

The overall objective will be to promote AI applications of high scientific merit among an extensive group of competent users, consistent with maintaining a responsive and productive computing environment. The initially approved loading is estimated to be about 30% of the total SUMEX-AIM capacity, leaving a substantial capacity for new research groups to enter at this time.

TECHNICAL CAPABILITY

The computer facility, consisting of a Digital Equipment Corporation model KI-10 central processor operating under the TENEX time-sharing monitor, is scheduled to be available in a limited fashion in June 1974, and to be fully operational in September. The system will have initially 197K words (36 bit) of high speed memory; 1.3M words of swapping storage; 40M words of disk storage; 2 9-track, 800 bpi industry-compatible tape units; 1 dual DEC tape unit; a line printer; and communications network interfaces providing user terminal access. At the present time the choice of a communications network approach is being finalized. However, specialized users will in any case find that the communications costs are small compared to the value of access to the system. This hardware complement may be expanded in the future based on available funding and justified user community needs.

Software support will evolve from the basic system as dictated by user research goals and requirements. Initially available programs will include a variety of TENEX user, utility, and text editor programs. Major user languages will include INTERLISP, SNOBOL, SAIL, FORTRAN-10, BLISS-10, BASIC, and MACRO-

Dr. Joshua Lederberg
SUMEX-AIM Computer Project
Stanford University Medical Center
Stanford, California 94305

(415) 497-5801